

Data Fusion of Geographically Dispersed Information: Experience With the Scalable Data Grid

Ke-Thia Yao, Ph.D., Craig E. Ward, and Dan M. Davis
Information Sciences Institute, USC, Marina del Rey, California

Test and evaluation (T&E) professionals today often face the distribution of data they need to use over significant distances. Some of these data sources are too large for easy transmission because of costs, delays, losses, security, and administrative burdens. The Information Sciences Institute (ISI) has been working with the Joint Forces Command on its transcontinentally distributed battlespace simulations, and they have conceived, architected, implemented, and tested a system that uses the data in place. This has been characterized as a scalable data grid that uses data cubes for rapid and focused retrieval. It is necessary that these data be available in a timely manner, organized for ease of access, securely stored, and easily manipulated for data mining. Suggestions are offered listing T&E situations to which the scalable data grid approach would seem applicable.

Key words: Aggregation and summarization; Apache Hadoop distributed files system; data collation; data mining; dimensions of interest; expectation-maximization (EM) algorithm; information from logged data; k-means algorithm; measures of performance; Meshrouter application; Urban Resolve Phase I.

Test and evaluation (T&E), while once done via handwritten analyses, has increasingly come to rely on computer collection and manipulation of data from a myriad of sensors and sources. In common with most of the rest of the defense community, T&E can now collect more data than it can easily analyze. More computing power allows for increases in the breadth and depth of the information collected. Now the same computing power must assist in identifying, ordering, storing, and providing easy access to that data. Fast networking allows large clusters of high performance computing resources, often distributed transcontinentally, to be brought to bear on T&E. This increase in fidelity has correspondingly increased the volumes of data that tests are capable of generating.

Coordinating distant computing resources and making sense of this mass of data is a problem that must be addressed. Unless data are analyzed and converted into information, testing will provide only a fraction of the knowledge that is possible. For the U.S. Joint Forces Command (USJFCOM) Urban Resolve exercises, which are used to evaluate new systems and sensors, Information Sciences Institute (ISI) developed a distributed logging system to capture publish-and-

subscribe messages from the high-level architecture (HLA) simulation federation. For a 2-week exercise, omitting nonessential data, this system logged over a terabyte of data (Yao and Wagenbreth, 2005).

In addition to the scalable data grid approach, the ISI team found that Hadoop provided a scalable, but conceptually simple, distributed computation paradigm that is based on map-reduce operations implemented over a highly parallel, distributed file system. Map-reduce implementations of k-means and expectation-maximization data mining algorithms were developed to take advantage of the Hadoop framework. This file system dramatically reduced the disk scan time needed by the iterative data mining algorithms. It was found that these algorithms could be effectively implemented across multiple Linux clusters connected over reserved high-speed networks. The data transmission reductions observed should be applicable in most T&E situations, even those that use lower bandwidth communications.

For this analysis, Hadoop jobs were created to experiment with the data mining performance characteristics in an environment that was based on connections to sites across widely dispersed geographic regions. Specially configured Linux cluster computers were installed at ISI in California, at the University of

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE MAR 2011		2. REPORT TYPE		3. DATES COVERED 00-00-2011 to 00-00-2011	
4. TITLE AND SUBTITLE Data Fusion of Geographically Dispersed Information: Experience With the Scalable Data Grid				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Information Sciences Institute, University of Southern California ,4676 Admiralty Way, Suite 1001, Marina del Rey, CA, 90292				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Illinois–Chicago (UIC) in Illinois, and at ISI-East in Virginia. All of these machines had large disk storage configurations, were located on network circuits capable of 10-Gb/s transmission, and remained dedicated to this research. The machine at ISI in California served as a control. Special network connectivity was established between UIC and ISI-East to test Hadoop across that geographic distance.

Agile data framework

In general, T&E analysts and customers have been interested not only in system data, but also in how well higher level mission tasks and objectives are satisfied. A measure of effectiveness (MOE) is a question, or a measure, designed to illuminate how well particular mission tasks are satisfied with respect to a system (Gentner et al. 1996).

A measure of performance is typically a quantitative measure of a system characteristic used to support an MOE. For example, sample MOE questions are “Can the enemy be located?” or “Can the system effectively detect movement within urban environment?” Measures of performance supporting these MOEs are typically statistical in nature.

Data loggers do extremely well at capturing detailed operational data. ISI’s distributed data loggers have captured terabytes of experimental data for the Urban Resolve Phase I exercises at JFCOM (Wagenbreth et al. 2010). Operational data included individual entity state changes and interactions among the entities. Depending on the type of entity, entity state changes may include location, orientation, and velocity. For vehicles, additional attributes may include external lights-on and engine-on. Interactions may include collision, damage assessment, sensor detection, and contact report.

The logged data collected from the test is often at too low a level to be of direct use to the T&E analysts. Information needs to be abstracted from the logged data by collation, aggregation, and summarization. To perform this data transformation, we had to define an analysis data model that is suitable for analysts and decision makers. ISI used a multidimensional data model as a way of representing the information from their perspective. Next, a logging data model representing the collected data has to be defined.

Then, to bridge the gap and connect two data models, ISI defined an abstraction relationship that mapped the logging model to the analysis model. This was a part of the scalable data grid toolkit (Yao and Wagenbreth, 2005).

ISI developed a sensor–target scoreboard that provided a visual way of quickly comparing the relative effectiveness of individual sensor platforms and sensor

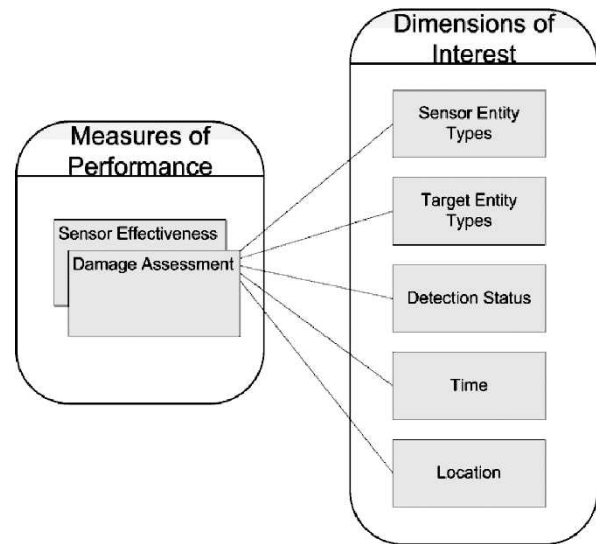


Figure 1. Dimensions of interest and measures of performance.

modes against different types of targets (Graebener et al. 2003, 2004). The sensor–target scoreboard was a specific instance of the more general multidimensional analysis (Kimbal et al. 1998). Such a scoreboard is an example of two dimensions of a multidimensional cube. Its two dimensions can be the sensor dimension and target. One can imagine extending the scoreboard to take into account, say, weather conditions and time of day. This would add two more dimensions to form a four-dimensional cube.

The analysis data model consists of two key concepts: dimensions of interest and measures of performance. Dimensions are used to define the coordinates of multidimensional data cubes. The cells within this data cube are the measure values (Figure 1).

Dimensions of interest

For large simulations, like the Urban Resolve exercises, the magnitude of data collected ranges into the terabytes. Dimensions categorize and partition the data along lines of interest to the analysts. Defining multiple crosscutting dimensions aids in breaking the data into smaller, orthogonal subsets.

Dimensions have associated measurement units, or coordinates. Choosing the granularity of these units aids in determining the size of the subsets. For example, depending on the dynamic nature of the phenomenon that the analysts are trying to study, they may choose to define the time dimension units in terms of minutes, days, weeks, or years.

Another dimension example is in terms of simulation entity groups. For the sensor dimension in the sensor–target scoreboard, the analyst may want to group together sensors by the type of platform: high-

flying unmanned aerial vehicles, midaltitude unmanned aerial vehicles, organic air vehicles, and unattended ground sensors. The targets may be grouped together, for example, by transportation mode: air, ground, and sea.

Hierarchical dimensions define how a coordinate relates to other coordinates as its subset. It serves to group together similar units from the analyst's perspective. By defining hierarchical dimensions, analysts inform the system about how to aggregate and summarize information. For example, the analysts may want to subdivide the sensor platform category into the sensor modes: moving target indicators, synthetic aperture radar, images, video, and acoustic.

Hierarchical dimensions can be viewed as a partial ordering relationship. At the top node of the partial ordering is the set containing all the coordinates. The bottom nodes of the partial ordering are singleton sets containing just one coordinate. Edges between nodes indicate a superset-subset relationship. A node's parent is its superset. A node's child is its subset. Nodes in hierarchical dimensions are also given coordinates. Calling the coordinates for nonsingleton nodes "abstract coordinates" is the accepted practice. Coordinates for singleton nodes are the concrete coordinate of the single element in the set.

Multiple decompositions of the same dimension are also useful. For example, there may be many different simulation federate types playing in the federation. The analysts may want to verify how each federate responds to the sensor contacts, so they can define the category to be the type of federates in which they are interested.

Measures of performance

After the data have been partitioned along lines of interest, the data subsets may still be large. Measures provide quantitative ways of characterizing the data subsets. A key characteristic that measures should have is that they can be aggregated. The hierarchical crosscutting dimensions partition the data into a hierarchy of subsets. The measure must be able to provide a meaningful summarization. To be computationally efficient, the measure aggregation operator must satisfy the associative and commutative properties—the measure of a set must be computable from the measures of its subsets. In the case of the sensor-target scoreboard, the measure of performance is simply an integer count of the number of times a sensor has detected a target. The aggregation operator is the addition operator.

RTI-s, a highly scalable implementation of the Runtime Infrastructure (RTI) (Helfinstine, Torpey, and Wagenbreth 2003), provides an interceptor plug-in framework that exposes calls to this HLA interface

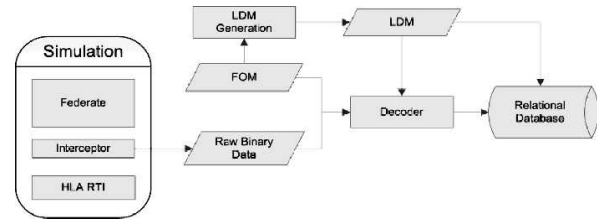


Figure 2. Logging data flow.

to registered plug-ins (Figure 2). The SDG exploits this plug-in to intercept and log messages and federate attribute updates and interaction sends. With respect to the RTI, the contents of these messages are raw binary strings. RTI provides publish-and-subscribe facilities to exchange these messages but not to decode their contents. To provide the query and analysis capabilities, these messages must be decoded with respect to the federation object model, the simulation definitional statement.

Cube computation

The data are made up of the facts or observations from a test. The facts have to be aggregated according to how hierarchical dimensions are defined. The cells that correspond to the abstract coordinates of the cube have to be computed. Here it is assumed that the aggregation operator satisfies the associative and commutative properties. Given these assumptions, measures can be efficiently computed for all the abstract coordinates by doing a bottom-up traversal of the partial ordering hierarchy.

T&E use of Hadoop

Easily obtained by anyone, Hadoop is an open source system, hosted by the Apache Software Foundation (Apache 2007). It provides a reliable, fault tolerant, distributed file system and application programming interfaces. These enable its map-reduce framework for the parallel evaluation of large volumes of test data.

The simplicity of the Hadoop programming model allows for straightforward implementations of many evaluation functions. The popular Java applications have the most direct access, but Hadoop also has streaming capabilities that allow for implementations in any preferred language.

Several other communities that need to handle large amounts of data are using map-reduce implementations to manage that data. Google started using a map-reduce system internally before 2004 (Dean and Ghemawat 2004). Yahoo runs the largest Hadoop cluster, running over a Linux cluster of over 10,000 cores (Yahoo 2008). Many vendors, e.g., Amazon, use

Hadoop as part of their cloud computing service. A list of organizations that make use of Hadoop can be found at the Hadoop wiki (Powered By, 2009).

In the 2008 terabyte sort challenge, Yahoo won by using Hadoop to sort 1 terabyte of data in 209 seconds (O'Malley 2008). That cluster consisted of 910 nodes with two quad core 2-GHz Xeons and 4 SATA disks per node.

Hadoop distributed file system

The Hadoop distributed files system (HDFS) runs on top of a native file system and is only accessible through the Hadoop application programming interfaces (APIs). HDFS configurations distribute data in equal-sized chunks across the available data nodes. This division of data works best for large files that can be stored as multiples of the chunking size configured for the HDFS. If the files are smaller than the chunking size, the HDFS will waste local file system resources with empty, allocated bytes. Redundancy and fault tolerance are achieved by replicating these chunks on multiple nodes. Hadoop attempts to run the map operations on copies of the data local to the mapping task. This reduces the amount of data that needs to be moved around.

ISI experiments used varying HDFS configurations. One configuration kept all nodes within a single rack. Another divided the nodes across half of the continental United States. Hadoop has three different node types: nodes for processing tasks, nodes for storing data, and a single node, called the name node, to coordinate the others. The tasks that are assigned to processing nodes are monitored for status. If a task appears to fail, it can be reassigned to another processing node. The assignments attempt to keep processing and data near each other, limiting the strain on any underlying communications resources, such as a network.

Distributed data mining algorithms

There are some T&E settings in which discovering unanticipated or novel data would be useful (Gehrig, Holloway, and Schroeter 2004). Data mining is a way of finding patterns in what otherwise would be random data. Many data mining algorithms are iterative in nature. They require the data to be scanned several times during the mining process. These algorithms can become prohibitively expensive for very large data sets that do not fit into memory and have to be stored on disk. Sequential disk access on a single disk can be several orders of magnitude slower than memory access. Hadoop with its potential to access thousands of disks in parallel provides a way of addressing this problem.

In addition, in some situations the test data themselves are distributed. For example, for JFCOM's Urban Resolve exercises, ISI implemented a distributed logger that stored high-level architecture runtime infrastructure (HLA RTI) messages locally, in situ where the messages were generated (Graebener et al. 2003; Yao and Wagenbreth 2005). Using Hadoop provides a convenient way to process the data without having to move them to a centralized location.

Two clustering algorithms

To test the feasibility of this approach, the ISI team implemented two data mining clustering algorithms in Hadoop: k-means and expectation-maximization (EM).

K-means is a popular data mining clustering algorithm that assigns a set of data instances into clusters (or subsets) based on some similarity metric. The k-means algorithm requires three inputs: an integer k to indicate the number of desired clusters as output, a distance function over the data instances, and the set of n data instances to be clustered. The distance of a data instance to itself is zero. The greater the distance between two data instances, the less similar the instances are. Typically, a data instance is represented as a numerical *vector*. The output of the algorithm is a set of k points representing the mean (or the center) of the k clusters. Each of the n data instances is assigned to the nearest cluster mean based on the distance function.

Below is pseudocode for the k-means algorithm:

1. Generate an initial guess for the k cluster (for example, by randomly selecting k points from the data instances as the k means).
2. Assign each of the n data instances to the nearest cluster mean.
3. Based on the data instance assignment, compute the new cluster mean for each of the k clusters.
4. While not done, go to Step 2.

Figure 3 illustrates some results of k-means clustering, correctly finding the means of the three distinct clusters. That is, given a set of points generated for this data set, the algorithm correctly discovered the patterns in the points.

The EM algorithm can be viewed as a probabilistic generalization of the k-means algorithm. Instead of representing a cluster by just its mean, EM represents a cluster by its mean and its variance (or covariance matrix), i.e., each cluster is represented by a Gaussian distribution. In addition, each cluster is associated with a weight, representing the probability of selecting the cluster. The sum of these k cluster weights is equal to 1.

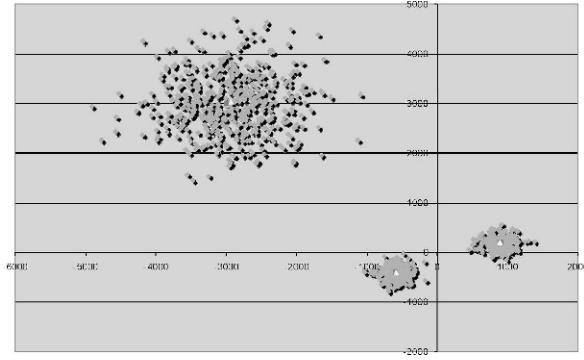


Figure 3. K-Means clustering of three distinct clusters of points.

This representation is called a Gaussian mixture model.

The steps of the EM algorithm are similar to the k-means algorithm. In Step 1 the initial guess now includes the k means, k variances, and k cluster weights. The assignment in Step 2, also known as the expectation step, is now slightly more complicated. Instead of assigning each data point to one cluster, each data point is assigned to each cluster with a probability based on a Gaussian distribution. In Step 3, the maximization step, the k means, k variances, and k cluster weights are recomputed based on the probabilistic assignment from Step 2.

Hadoop implementation

Only the Hadoop implementation of the k-means algorithm will be described, the EM Hadoop implementation being similar. There exists a variety of ways to generate the initial guess in Step 1. If there is a priori knowledge of the range of values of the data, then k means can be generated randomly using a uniform distribution. Otherwise, scanning the data instances once will allow computing the range values. Or, scanning the data instances and randomly selecting k instances as the means is possible. To simplify the algorithm description, we shall assume that there is a priori knowledge.

Step 2 corresponds to the map operation. Map functions have the form

Map: $(in-key, in-value) \rightarrow list (out-key, out-value)$

In this case, the $in-key$ is null, and the $in-value$ is the data instance vector. The $out-key$ is an integer from 1 to k representing the cluster identifier, and the $out-value$ is a list of pairs, where each pair consists of the data instance vector and the integer 1.

K-means map: $(null, data-instance) \rightarrow list (cluster-id, (data-instance, 1))$

Step 3 corresponds to the reduce operation. Reduce functions have the form:

Reduce: $(in-key, list (in-value)) \rightarrow list (out-key, out-value)$

In this case the input $(in-key, in-value)$ is the output of the k-means map $(cluster-id, (data-instance, 1))$. For each $cluster-id$, the reduce operation sums all the $(data-instance, 1)$ pairs associated with that cluster-id.

K-means reduce: $(cluster-id, (data-instance, count)) \rightarrow list (cluster-id, (sum-of- data-instances, number-of-instances))$

Here the $sum-of- data-instances$ divided by $number-of-instances$ is the mean of the cluster.

Here is a simple, but naïve, Hadoop implementation of the k-means algorithm:

1. Random generate k points as initial k means.
2. Apply k-means map and reduce.
3. While not done, go to Step 2.

Test environment setup

Data mining Hadoop jobs were created for the SIMC-IC project to experiment with the performance characteristics of Hadoop in an environment that provided high-speed network connections to sites across large geographic regions. As mentioned before, clusters in California, Illinois, and Virginia were connected via a high-bandwidth link.

Each cluster machine was composed of

- 10 nodes
- 5.3-TB local disk
- 2 clusters running Fedora 8
- 1 cluster running Debian
- 1 10-G E network card
- 1 1-G card for management only
- Dual quad core (8 cores per node) CPUs

The version of Hadoop used for the experiments was 0.17.2.1. Each cluster used the Java SE Runtime Environment 1.6 (build 1.6.0_11-b03).

Hadoop clusters were configured using the available nodes such that both the control Hadoop cluster and the distributed Hadoop cluster had the same number of nodes, one name node and nine nodes running data and job task services. The only difference was that the control cluster used only local network connections while the other used wide area network connections.

For the wide area network Hadoop cluster, two configurations were used. One configuration used the default network resources and one used dedicated Internet 2 high-bandwidth lines reserved for short periods.

Table 1. Data load test results.

	User	System	Elapsed
ISI Local	44.85	22.09	2:05.69
ISIE/UIC (standard)	46.98	18.38	14:27.75
ISIE/UIC (fastnet)	49.18	18.94	29:20.78

Data load

In addition to the data mining jobs developed, the ability of Hadoop to load and store data was tested. A simple data load of six 1.2-GB files was performed using the default settings, each block of data replicated on three nodes. All time data was collected from the *time(1)* command.

As would be expected, the quickest data loads were with the local nodes configuration (Table 1). The actual processing times were not that much different for each configuration. The major difference was in clock time, indicating that the distributed systems spent significant time in suspended wait states while the network subsystems performed their functions. The Fastnet version using Internet 2 actually took longer elapsed time than the standard version. However, during the execution of the Fastnet version, we observed Java network exceptions being thrown.

Data mining jobs

Two implementations of the k-means algorithm were used to test the processing capabilities of Hadoop. An EM job was also developed, but this job was not used for this experiment. The UIC Angle data set was searched for points within the data where the data clustered. One implementation used a “naive” approach while the other used a more efficient “smart” approach. The naive implementation did not use the combine step allowed by the Hadoop API. This resulted in much more network usage because more data had to be passed among the task nodes. The smart implementation made use of this step and greatly reduced the amount of data exchanged.

The k-means jobs iterated over the data set with an initial set of cluster points, each time updating the set of cluster points to better fit the data, each resulting set of cluster points becoming the input for the next iteration. When either the points stopped significantly changing or the maximum number of iterations was reached, the job stopped.

For development and initial testing, the job was tested using points randomly generated using known center coordinates. The results of a run were expected to match the input provided to the random point generator (Table 2).

As with the data loads, the data mining jobs performed best on the local nodes setup. The

Table 2. K-means results.

	User	System	Elapsed
ISI Local (smart)	1.68	0.18	1:37.76
ISI Local (naive)	6.55	0.92	40:38.64
ISIE/UIC (smart/stand)	1.67	0.19	1:52.80
ISIE/UIC (smart/Fastnet)	2.25	0.27	8:25.08
ISIE/UIC (naive/stand)	5.35	0.96	1:12:03
ISIE/UIC (naive/Fastnet)	8.40	1.72	2:14:16 KILLED*

*The naive run was killed at the elapsed time in the seventh job iteration. The maximum number for a run is 32.

differences between local and networked systems are not as pronounced as with the data loads. This is likely due to the ability of Hadoop to process chunks of data in a “rack-aware” manner. The smart implementations tended to not require long haul network services and were able to process data in what to them was a local manner. Again, the Fastnet version took longer elapsed time than the standard version.

Network utilization

In the previous section, our experiments exercised Hadoop across differing network configurations. One configuration used the “normal” connectivity found in the network while another ran Hadoop over special high-speed links with a theoretical peak throughput of 10 Gb/s. But, Hadoop results did not reflect the advantage of the high-speed links.

To rule out the possibility that the high-speed links were faulty, we used another software system to obtain independent measurements. The tool used to test this capability was the Meshrouter, which was designed for high throughput HLA RTI communications (Barrett and Gottschalk 2004; Brunett and Gottschalk 1998). The tests show the Meshrouter application is capable of achieving 1.5 Gb/s with a single TCP stream, and up to 5 Gb/s with combined streams.

Based on this throughput experiment, we reasoned that Hadoop was not able to take full advantage of the high-speed network. As mentioned previously, it was observed that Java generated network exceptions during the execution. Although Hadoop is designed to be fault tolerant, the exceptions most likely slowed down its execution.

Moreover, to achieve 50 percent capacity of the high-speed network, the Meshrouter application required several TCP streams. It is suspected that even without the network exceptions, Hadoop will not be able to take full advantage of the 10-Gb/s network.

Below, the details of the high-speed network throughput experiment using the Meshrouter are described. The Meshrouter and associated applications implement interest managed communication (RTI) used by several entity simulators in general use. Test

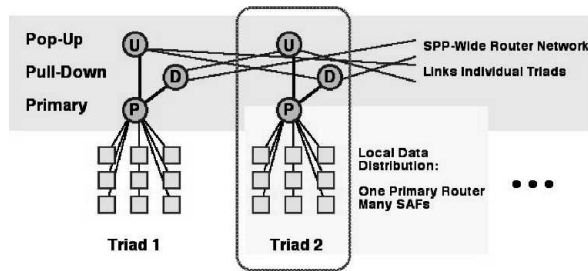


Figure 4. Schematic meshrouter topology.

programs named publish and subscribe were used to exercise the network in a controlled and repeatable manner. The Meshrouter is a complex real-world application.

The bandwidth experiments were conducted using the standard ISI Meshrouter formalism for interest-managed communications. A schematic of the Meshrouter is shown in Figure 4.

The overall communications scheme consists of collections of processors (labeled “SAFs” in this legacy diagram) each communicating with a specified “primary” router (P). Interest-limited message exchange among the various basic “triads” is done using a network of additional “pop-up” and “pull-down” routers. As is described in Barrett and Gottschalk (2004), the three routers on a triad are instanced as separate objects within a single Meshrouter process.

The execution of an actual message transfer is implemented by a software stack as shown in Figure 5.

The results reported here use an RTI-s implementation for both interest enumeration and the lowest-level communications primitives (“data flow nodes”). While this has enormous advantages, it does have the generic disadvantage of any general purpose “plug and play” system in terms of significant, incompletely understood overheads.

Standard RTI-s data flow implementations exist for both TCP and UDP communications. The results presented here use the TCP implementation.

The application processes for the benchmark tests are of two forms:

Publish Processors: Send out messages of specified length and interest state. The nominal total publication rate (megabyte per second) is controlled by a data file that is reread periodically (by all publish processors). This means that the nominal experimental data rate can be controlled dynamically.

Subscribe Processors: Receive messages for a specified interest state, collecting messages from multiple publishers, as appropriate. The subscribe processes are instrumented to measure actual

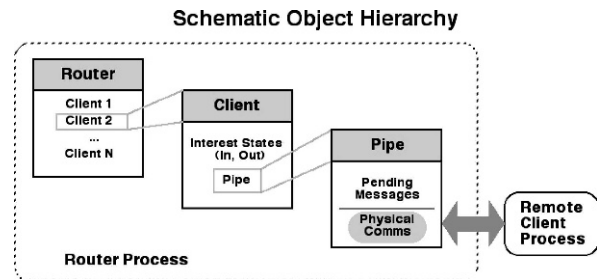


Figure 5. Factored Meshrouter implementation, with application-specific communications primitives.

incoming message rates and to detect missed messages.

The routers in Figure 5 direct individual messages from publishers to subscribers according to the interest declarations. The router processes are also instrumented to determine the fraction of (wall clock) time spent in communications management (vs. simply waiting for input).

Two modes were tested. In the first mode, a single TCP connection was set up between a pair of Meshrouters at distant locations. The measured bandwidth was approximately 300 Mb/s. The second mode used eight mesh routers at each site, each with multiple clients and multiple TCP connections. Measured aggregate bandwidth was approximately 4.6 Gb/s.

This test demonstrated that 50 percent of the capacity of the high speed wide area network can be effectively employed by a real world application.

Programming Hadoop

For the T&E programmers, Hadoop should be an easy system to use. Installation was straightforward. The rapid changes in Hadoop releases made keeping up problematic; some releases broke existing code. In the middle of a test evolution, one might be well advised not to install every update, but “freeze” Hadoop for the duration of the test.

Shell scripts might be found useful to reduce the complexity of setup, change, and maintenance of the various Hadoop configurations across sites. ISI experience was that once these configurations were in place, changing them was an easy and straightforward operation.

Development was convenient for Hadoop jobs. Running and debugging standalone Hadoop jobs in the Eclipse IDE allowed rapid turnaround on application bug fixes.

Conclusions

This article supports the proposition that the implementation and use of the SDG and Hadoop

show promise for the T&E environments. It reported on experiments implementing the SDG and on using distributed data analysis and data mining implemented over the Apache Hadoop framework. ISI's experience is that the SDG and Hadoop provided a scalable, but conceptually simple, distributed computation paradigm based on the standard map and reduce operations implemented over a highly parallel, distributed file system. ISI found it practical to develop map and reduce implementations of k-means and EM data mining algorithms that took advantage of the Hadoop framework. The Hadoop file system dramatically reduced the disk scan time needed by these iterative data mining algorithms. ISI has successfully executed these algorithms across multiple Linux clusters over dedicated 10-Gb/s networks. The ISI team holds that the results of these experiments support the potential for the use of these tools in T&E. □

DR. KE-THIA YAO is a project leader and research scientist at the University of Southern California Information Sciences Institute. His research has been centered on the JESPP project, which has the goal of supporting very large-scale distributed military simulations involving millions of autonomous agent entities. He has developed a suite of monitoring/logging/analysis tools to help users better understand the computational and behavioral properties of large-scale simulations and has designed the Scalable Data Grid for distributed data management of large archives. He received his bachelor of science degree in EECS from the University of California, Berkeley, and his master of science and doctor of philosophy degrees in computer science from Rutgers University, New Brunswick, New Jersey. E-mail: kyao@isi.edu

CRAIG E. WARD is a parallel computer systems analyst at the Information Sciences Institute. Much of his recent research has focused on large-scale data management in the defense and the health sectors. His concentration has been on open source tools. Previously, he performed computer analysis for law enforcement in California. He has a bachelor of arts in history from the University of California, Irvine, and a master of science degree in computer science from Loyola Marymount University in Los Angeles, California. E-mail: cward@isi.edu

DAN M. DAVIS is the director, JESPP Project, Information Sciences Institute (ISI), University of Southern California (USC), focusing on large-scale distributed DoD simulations. As the assistant director of the Center for Advanced Computing Research at Caltech, he managed Synthetic Forces Express, bringing HPC to DoD simulations. Prior experience includes work as a software engineer at the Jet Propulsion Laboratory and at a classified

project at Martin Marietta. He saw duty in Vietnam as a cryptologist in the USMC, and he retired as a commander, cryptologic specialty, USNR. He received a bachelor of arts and a juris doctor degree in law from the University of Colorado in Boulder. E-mail: ddavis@isi.edu

References

- Apache. 2007. *The Hadoop distributed file system: Architecture and design*. http://hadoop.apache.org/core/docs/r0.17.1/hdfs_design.pdf (accessed June 25, 2009).
- Barrett, B., and T. Gottschalk. 2004. Advanced message routing for scalable distributed simulations. In *Proceedings of the Interservice/Industry Training, Simulation and Education Conference*, Orlando, Florida.
- Brunett, S., and T. Gottschalk. 1998. A large-scale metacomputing framework for the ModSAF real-time simulation. *Parallel Computing* 24: 1873–1900.
- Dean, J., and S. Ghemawat. 2004. *MapReduce: Simplified data processing on large clusters*. operating system design and implementation. San Francisco: Operating System Design and Implementation. <http://labs.google.com/papers/mapreduce.html> (accessed June 25, 2009).
- Gehrig, J. F., G. Holloway, and G. Schroeter. 2002. Reflections on test and evaluation: T&E infrastructure, reengineering Army T&E, and building a viable test range complex. *Program Manager*. July–August, 2002.
- Genter, F. C., P. S. Best, and P. H. Cunningham. 1999. *Sources of Measures of Effectiveness (MOEs) for Assessing Human Performance in Aeronautical Systems*, Pensacola, Florida: Navy Advancement Center.
- Graebener, R., G. Rafuse, R. Miller, and K.-T. Yao. 2003. The road to successful joint experimentation starts at the data collection trail. In *Proceedings of the Interservice/Industry Training, Simulation and Education Conference*. July–Aug 2002, Orlando, Florida, 56–62.
- Graebener, R., G. Rafuse, R. Miller, and K.-T. Yao. 2004. The road to successful joint experimentation starts at the data collection trail—Part II. In *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference*. Orlando, Florida.
- Helfinstine, B., M. Torpey, and G. Wagenbreth. 2003. Experimental interest management architecture for DCEE. In *Interservice/Industry Training, Simulation, and Education Conference*. Orlando, Florida.
- Kimbal, R., L. Reeves, M. Ross, and W. Thornwaite. 1998. *The Data Warehouse Lifecycle Toolkit*. Hoboken, New Jersey: Wiley.
- O'Malley, O. 2008. *TeraByte sort on Apache Hadoop*. <http://www.hpl.hp.com/hosted/sortbenchmark/YahooHadoop.pdf>. (accessed June 29, 2009).

Rathnam, T., and C. J. J. Paredis. 2004. Developing federation object models using ontologies. In *Proceedings of the 2004 Winter Simulation Conference*, Washington, DC.

Wagenbreth, G., G. E. Ward, K.-T. Yao, and D. M. Davis. 2010 (Pending). Non-disruptive data logging: Tools for JFCOM large-scale simulations. In *Proceedings of the Simulation Interoperability Workshop*. Orlando, Florida.

Yahoo. 2008. *Yahoo! launches world's largest Hadoop production application*. <http://developer.yahoo.net/blogs/hadoop/2008/02/yahoo-worlds-largest-production-hadoop.html> (accessed June 29, 2009).

Yao, K. T., and G. Wagenbreth. 2005. Simulation data grid: joint experimentation data management and analysis. In *Proceedings of the Interservice/Industry Training, Simulation and Education Conference*. Orlando, Florida.

Acknowledgments

We would like to thank our colleagues at JFCOM, and at UIC, especially Professor Robert Grossman and the Messers Michael Sabala and Yuhong Gu. Further, this work would not have been possible without the assistance of Tom Lehman of ISI East. This material is based in part on research sponsored by the Air Force Research Laboratory under agreement number FA8750-05-2-0204. The U.S. government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory or the U.S. government.

The University of Alabama in Huntsville

Division of Professional and
Continuing Studies

Gain knowledge and skills throughout the entire system life cycle in two specialty engineering disciplines with emphasis on Aerospace and Defense. Take one course or the entire certificate program. Both certificate programs are offered online, on demand – allowing you to begin a course anytime and complete in one month. Classroom courses are available on the UAHuntsville campus or at your location.

TEST & EVALUATION

CERTIFICATE

Experts lead you through the T&E process of the first requirements dialogue between systems and T&E personnel to a requirement-responsive data report. Learn the roles, responsibilities, and functions applied across the life cycle using a multi-disciplinary approach.

CORE COURSES:

Test and Evaluation Foundation
Requirements Development
Data Collection, Analysis and Reporting
Test Design, Planning and Conduct

ELECTIVES (choose one):

System Validation and Verification
Modeling and Simulation for T&E

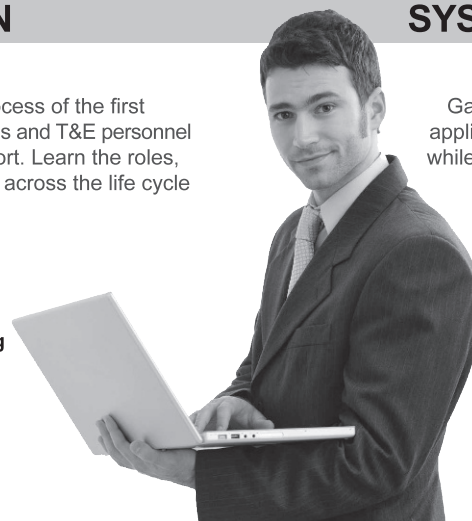
SYSTEMS ENGINEERING

CERTIFICATE

Gain an understanding of the process and application of systems engineering methods while refining your ability to design, develop, and manage complex systems.

Systems Engineering Overview

Requirements Development
Trade Studies and Decision Making
Risk Assessment and Management
System Validation and Verification
Systems Analysis, Modeling,
and Simulation



An Affirmative Action/Equal Opportunity Institution

Visit www.PCS.uah.edu/ITEA for details • 256.824.4238

Available Classroom and Online